

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

REALTIME DATA LLC d/b/a IXO,

Plaintiff,

v.

MONGODB INC.,

Defendant.

C.A. No. _____

JURY TRIAL DEMANDED

**COMPLAINT FOR PATENT INFRINGEMENT
AGAINST MONGODB, INC.**

This is an action for patent infringement arising under the Patent Laws of the United States of America, 35 U.S.C. § 1 *et seq.* in which Plaintiff Realtime Data LLC d/b/a IXO (“Plaintiff,” “Realtime,” or “IXO”) makes the following allegations against Defendant MongoDB Inc., (“MongoDB” or “Defendant”):

PARTIES

1. Realtime is a limited liability company organized under the laws of the State of New York. Realtime has its principle place of business at 66 Palmer Avenue, Suite 27, Bronxville, NY 10708. Since the 1990s, Realtime has researched and developed specific solutions for data compression, including, for example, those that increase the speeds at which data can be stored and accessed. As recognition of its innovations rooted in this technological field, Realtime holds 40 United States patents and has numerous pending patent applications. Realtime has licensed patents in this portfolio to many of the world’s leading technology companies. The patents-in-suit relate to Realtime’s development of advanced systems and methods for fast and efficient data compression using numerous innovative compression techniques based on, for example, particular attributes of the data.

2. On information and belief, MongoDB is a Delaware corporation with its principal place of business at 1633 Broadway 38th Floor, New York, NY 10019. MongoDB can be served through its registered agent, Corporation Service Company, 251 Little Falls drive, Wilmington, Delaware, 19808.

JURISDICTION AND VENUE

3. This action arises under the patent laws of the United States, Title 35 of the United States Code. This Court has original subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

4. This Court has personal jurisdiction over Defendant MongoDB in this action because MongoDB is incorporated in Delaware and has committed acts within the District of Delaware giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over MongoDB would not offend traditional notions of fair play and substantial justice. MongoDB, directly and through subsidiaries or intermediaries, has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents.

5. Venue is proper in this district under 28 U.S.C. § 1400(b). Upon information and belief, MongoDB is incorporated in Delaware, has transacted business in the District of Delaware, and has committed acts of direct and indirect infringement in this District.

COUNT I **INFRINGEMENT OF U.S. PATENT NO. 9,116,908**

6. Plaintiff realleges and incorporates by reference the foregoing paragraphs, as if fully set forth herein. Plaintiff Realtime is the owner by assignment of United States

Patent No. 9,116,908 (the “’908 Patent”) entitled “System and methods for accelerated data storage and retrieval.” The ’908 Patent was duly and legally issued by the United States Patent and Trademark Office on August 25, 2015. A true and correct copy of the ’908 Patent is included as Exhibit A.

7. On information and belief, MongoDB has offered for sale, sold and/or imported into the United States MongoDB products and services that infringe the ’908 Patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, MongoDB products and services, *e.g.*, MongoDB 3.0 and above, MongoDB Ops Manager, MongoDB Cloud Manager, WiredTiger storage engine, and the system hardware on which they operate, and all versions and variations thereof since the issuance of the ’908 Patent (“Accused Instrumentalities”).

8. On information and belief, MongoDB has directly infringed and continues to infringe the ’908 Patent, for example, by making, selling, offering for sale, and/or importing the Accused Instrumentalities, and through its own use and testing of the Accused Instrumentalities, which constitute performing a method for accelerating data storage of data claimed by Claim 21 of the ’908 Patent, comprising: compressing a first data block with a first data compression technique to provide a first compressed data block; and compressing a second data block with a second data compression technique to provide a second compressed data block, wherein the first data compression technique and the second data compression technique are different; storing the first and second data compressed blocks on a memory device wherein the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form. Upon information and belief, MongoDB uses the Accused

Instrumentalities, which perform the infringing method, for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to its customers.

9. MongoDB also indirectly infringes the '908 Patent by manufacturing, using, selling, offering for sale, and/or importing the accused products, with knowledge that the accused products were and are especially manufactured and/or especially adapted for use in infringing the '908 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use. On information and belief, the Accused Instrumentality is designed to function with compatible hardware to perform a method for accelerating data storage of data comprising: compressing a first data block with a first data compression technique to provide a first compressed data block; and compressing a second data block with a second data compression technique to provide a second compressed data block, wherein the first data compression technique and the second data compression technique are different; storing the first and second data compressed blocks on a memory device wherein the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form. Because the Accused Instrumentality is designed to operate as the claimed method for accelerating data storage of data, the Accused Instrumentality has no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. MongoDB's manufacture, use, sale, offering for sale, and/or importation of the Accused Instrumentality constitutes contributory infringement of the '908 Patent.

10. On information and belief, MongoDB has had knowledge of the '908 Patent since at least the filing of the original Complaint in this action, or shortly thereafter, and

on information and belief, MongoDB knew of the '908 Patent and knew of its infringement, including by way of this lawsuit.

11. MongoDB's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe claims of the '908 Patent. Use of the Accused Instrumentalities in their ordinary and customary manner results in infringement of claims of the '908 Patent. For example, MongoDB explains to customers the benefits of using the Accused Instrumentalities, such as by touting their performance advantages: "[O]nly the differences of each successive snapshot are stored. Compression and block-level deduplication reduce the size of snapshot data." See <https://docs.opsmanager.mongodb.com/current/application/>. For similar reasons, MongoDB also induces its customers to use the Accused Instrumentalities to infringe other claims of the '908 Patent. MongoDB specifically intended and was aware that the normal and customary use of the Accused Instrumentalities on compatible systems would infringe the '908 Patent. MongoDB performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '908 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, MongoDB engaged in such inducement to promote the sales of the Accused Instrumentalities, *e.g.*, through MongoDB's user manuals, product support, marketing materials, and training materials to actively induce the users of the accused products to infringe the '908 Patent. Accordingly, MongoDB has induced and continues to induce end users of the accused products to use the accused products in their ordinary and customary

way with compatible systems to make and/or use systems infringing the '908 Patent, knowing that such use of the Accused Instrumentalities with compatible systems will result in infringement of the '908 Patent.

12. The Accused Instrumentalities compress a first data block with a first data compression technique to provide a first compressed data block. For example, the Accused Instrumentalities support Snappy and zlib compression techniques (e.g., “By default, collection data in WiredTiger use Snappy block compression; zlib compression is also available. See <https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>.) As another example, the Accused Instrumentalities provide MongoDB Ops Manager that implements a compression technique (e.g., “Compression and block-level deduplication reduce the size of snapshot data.” See <https://docs.opsmanager.mongodb.com/current/application/>.)

13. The Accused Instrumentalities compress a second data block with a second data compression technique to provide a second compressed data block, wherein the first data compression technique and the second data compression technique are different. For example, the Accused Instrumentalities include a data deduplication technique (e.g., “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.” <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>.) As such, the Accused Instrumentalities define the data deduplication technique as “[A] process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” See

<https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>.

As another example, the Accused Instrumentalities perform index prefix compression that “deduplicates common prefixes from indexed fields.” *See*

<https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>. Moreover, the

Accused Instrumentalities further specify that index prefix compression “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.

As another example, the Accused Instrumentalities support Snappy and zlib compression techniques (e.g., “By default, collection data in WiredTiger use Snappy block compression; zlib compression is also available.” *See*

<https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>.) As such,

Snappy is a LZ77-type compressor and zlib is an LZ77 variant called deflation. *See*

https://zlib.net/zlib_tech.html; https://github.com/google/snappy/blob/master/format_description.txt. In contrast to LZ77-type compression, data deduplication is a “process that

eliminates redundant data.” More specifically, the Accused Instrumentalities explain that

“data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning

or solid state disks.” *See*

<https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>

14. The Accused Instrumentalities store the first and second data compressed blocks on a memory device wherein the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed

form. For example, the Accused Instrumentalities store data snapshots “in a managed MongoDB database in a compressed, de-duplicated format.” *See* <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>. As another example, the Accused Instrumentalities provide routines that return the size of compressed data on disk storage (e.g., “db.collection.storageSize() to return the size in bytes of the collection on disk storage. If collection data is compressed (which is the default for WiredTiger), the storage size reflects the compressed size and may be smaller than the value returned by db.collection.dataSize().” *See* <https://docs.mongodb.com/manual/faq/storage/index.html>.) As another example, the Accused Instrumentalities state that index prefix deduplication compression “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See* <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>. Due to the data reduction and acceleration features of the specific compression algorithms used, the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. As such, in the Accused Instrumentalities “[W]iredTiger's granular concurrency control and native compression will provide the best all-around performance and storage efficiency for the broadest range of applications.” *See* Performance Best Practices for MongoDB.

15. MongoDB also infringes other claims of the '908 Patent, directly and through inducing infringement and contributory infringement.

16. On information and belief, use of the Accused Instrumentalities in their ordinary and customary fashion results in infringement of the methods claimed by the '908 Patent.

17. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' data storage accelerating features, MongoDB has injured Realtime and is liable to Realtime for infringement of the '908 Patent pursuant to 35 U.S.C. § 271.

18. As a result of MongoDB's infringement of the '908 Patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for MongoDB's infringement, but in no event less than a reasonable royalty for the use made of the invention by MongoDB, together with interest and costs as fixed by the Court.

COUNT II
INFRINGEMENT OF U.S. PATENT NO. 9,667,751

19. Plaintiff realleges and incorporates by reference the foregoing paragraphs, as if fully set forth herein.

20. Plaintiff Realtime is the owner by assignment of United States Patent No. 9,667,751 (the "'751 Patent") entitled "Data feed acceleration." The '751 Patent was duly and legally issued by the United States Patent and Trademark Office on May 30, 2017. A true and correct copy of the '751 Patent is included as Exhibit B.

21. On information and belief, MongoDB has offered for sale, sold and/or imported into the United States MongoDB products and services that infringe the '751 Patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, MongoDB products and services, *e.g.*, MongoDB 3.0 and above, MongoDB Ops Manager, MongoDB Cloud Manager, WiredTiger storage

engine, and the system hardware on which they operate, and all versions and variations thereof since the issuance of the '751 Patent ("Accused Instrumentalities").

22. On information and belief, MongoDB has directly infringed and continues to infringe the '751 Patent, for example, through its own use and testing of the Accused Instrumentalities, which in the ordinary course of their operation perform a method for compressing data claimed by Claim 1 of the '751 Patent, comprising: analyzing content of a data block to identify a parameter, attribute, or value of the data block that excludes analyzing based solely on reading a descriptor; selecting an encoder associated with the identified parameter, attribute, or value; compressing data in the data block with the selected encoder to produce a compressed data block, wherein the compressing includes utilizing a state machine; and storing the compressed data block; wherein the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. Upon information and belief, MongoDB uses the Accused Instrumentalities, which perform the infringing method, for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to MongoDB's customers.

23. On information and belief, MongoDB has had knowledge of the '751 Patent since at least the filing of the original Complaint in this action, or shortly thereafter, and on information and belief, MongoDB knew of the '751 Patent and knew of its infringement, including by way of this lawsuit.

24. Upon information and belief, MongoDB's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and

technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 1 of the '751 Patent by analyzing content of a data block to identify a parameter, attribute, or value of the data block that excludes analyzing based solely on reading a descriptor; selecting an encoder associated with the identified parameter, attribute, or value; compressing data in the data block with the selected encoder to produce a compressed data block, wherein the compressing includes utilizing a state machine; and storing the compressed data block; wherein the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. For example, MongoDB explains to customers the benefits of using the Accused Instrumentalities, such as by touting their efficiency: “[O]nly the differences of each successive snapshot are stored. Compression and block-level deduplication reduce the size of snapshot data.” <https://docs.opsmanager.mongodb.com/current/application/>. For similar reasons, MongoDB also induces its customers to use the Accused Instrumentalities to infringe other claims of the '751 Patent. MongoDB specifically intended and was aware that these normal and customary activities would infringe the '751 Patent. MongoDB performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '751 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, MongoDB engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, MongoDB has induced and continues to induce users of the accused products to use the accused products in their ordinary and

customary way to infringe the '751 Patent, knowing that such use constitutes infringement of the '751 Patent.

25. MongoDB also indirectly infringes the '751 Patent by manufacturing, using, selling, offering for sale, and/or importing the accused products, with knowledge that the accused products were and are especially manufactured and/or especially adapted for use in infringing the '751 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use. On information and belief, the Accused Instrumentality is designed to perform a method for compressing data comprising: analyzing content of a data block to identify a parameter, attribute, or value of the data block that excludes analyzing based solely on reading a descriptor; selecting an encoder associated with the identified parameter, attribute, or value; compressing data in the data block with the selected encoder to produce a compressed data block, wherein the compressing includes utilizing a state machine; and storing the compressed data block; wherein the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. Because the Accused Instrumentality is designed to operate as the claimed method for compressing, the Accused Instrumentality has no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. MongoDB's manufacture, use, sale, offering for sale, and/or importation of the Accused Instrumentality constitutes contributory infringement of the '751 Patent.

26. The Accused Instrumentalities analyze content of a data block to identify a parameter, attribute, or value of the data block that excludes analyzing based solely on reading a descriptor. For example, the Accused Instrumentalities support a data

deduplication technique (e.g., “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.” [https://docs.opsmanager.mongodb.com/current/core/administration-interface/.](https://docs.opsmanager.mongodb.com/current/core/administration-interface/)) As such, the Accused Instrumentalities define data deduplication as “[A] process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” See <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>. As another example, the Accused Instrumentalities perform index prefix compression that “deduplicates common prefixes from indexed fields.” See <https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>. Moreover, the Accused Instrumentalities further specify that index prefix compression “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.

27. The Accused Instrumentalities select an encoder associated with the identified parameter, attribute, or value. For example, the Accused Instrumentalities support data compression and deduplication techniques (e.g., “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.” [https://docs.opsmanager.mongodb.com/current/core/administration-interface/.](https://docs.opsmanager.mongodb.com/current/core/administration-interface/)) As another example, the Accused Instrumentalities support index prefix deduplication compression, which “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of

memory.” <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>. As such, the Accused Instrumentalities select data deduplication technique to eliminates redundant data. In the Accused Instrumentalities, “[O]nly unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>. The Accused Instrumentalities use data compression technique to address redundancy in unique documents or blocks.

28. The Accused Instrumentalities compress data in the data block with the selected encoder to produce a compressed data block, wherein the compressing includes utilizing a state machine. For example, the Accused Instrumentalities provide data compression and deduplication techniques outputting compressed data blocks (e.g., “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.” <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>.) In particular, in the Accused Instrumentalities, deduplication is a “process that eliminates redundant data.” *See*

<https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>. As such, the Accused Instrumentalities explain that during deduplication “[O]nly unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>.

As another example, the Accused Instrumentalities support index prefix deduplication compression, which “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.

compression. As another example, the Accused Instrumentalities perform Snappy or zlib compression (e.g., “By default, collection data in WiredTiger use Snappy block compression; zlib compression is also available. *See* <https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine.>)

29. The Accused Instrumentalities store the compressed data block. For example, the Accused Instrumentalities store data blocks “in a managed MongoDB database in a compressed, de-duplicated format.” *See* <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>. As another example, the Accused Instrumentalities provide routines that return the size of compressed data on disk storage (e.g., “db.collection.storageSize() to return the size in bytes of the collection on disk storage. If collection data is compressed (which is the default for WiredTiger), the storage size reflects the compressed size and may be smaller than the value returned by db.collection.dataSize().” *See* <https://docs.mongodb.com/manual/faq/storage/index.html.>) As another example, the Accused Instrumentalities state that index prefix deduplication compression “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See* <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression.>

30. The Accused Instrumentalities compress data, wherein the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. For example, the Accused Instrumentalities perform compression and data de-duplication (e.g., “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.”

[https://docs.opsmanager.mongodb.com/current/core/administration-interface/.](https://docs.opsmanager.mongodb.com/current/core/administration-interface/)) Due to the data reduction and acceleration features of the specific compression algorithms used, the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. As such, in the Accused Instrumentalities “[W]iredTiger’s granular concurrency control and native compression will provide the best all-around performance and storage efficiency for the broadest range of applications.” *See* Performance Best Practices for MongoDB.

31. On information and belief, MongoDB also infringes, directly and through induced infringement, and continues to infringe other claims of the ’751 Patent.

32. On information and belief, use of the Accused Instrumentalities in their ordinary and customary fashion results in infringement of the methods claimed by the ’751 Patent.

33. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, MongoDB has injured Realtime and is liable to Realtime for infringement of the ’751 Patent pursuant to 35 U.S.C. § 271.

34. As a result of MongoDB’s infringement of the ’751 Patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for MongoDB’s infringement, but in no event less than a reasonable royalty for the use made of the invention by MongoDB, together with interest and costs as fixed by the Court.

COUNT III
INFRINGEMENT OF U.S. PATENT NO. 8,933,825

35. Plaintiff realleges and incorporates by reference the foregoing paragraphs, as if fully set forth herein. Plaintiff Realtime is the owner by assignment of United States Patent No. 8,933,825 (the “’825 Patent”) entitled “Data compression systems and methods.” The ’825 Patent was duly and legally issued by the United States Patent and Trademark Office on January 13, 2015. A true and correct copy of the ’825 Patent is included as Exhibit C.

36. On information and belief, MongoDB has offered for sale, sold and/or imported into the United States MongoDB products and services that infringe the ’825 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, MongoDB products and services, *e.g.*, MongoDB 3.0 and above, MongoDB Ops Manager, MongoDB Cloud Manager, WiredTiger storage engine, and the system hardware on which they operate, and all versions and variations thereof since the issuance of the ’825 Patent (“Accused Instrumentalities”).

37. On information and belief, MongoDB has directly infringed and continues to infringe the ’825 Patent, for example, by making, selling, offering for sale, and/or importing the Accused Instrumentalities, and through its own use and testing of the Accused Instrumentalities, which constitute performing a method claimed by Claim 18 of the ’825 Patent, comprising: associating at least one encoder to each one of a plurality of parameters or attributes of data; analyzing data within a data block to determine whether a parameter or attribute of the data within the data block is identified for the data block; wherein the analyzing of the data within the data block to identify a parameter or attribute of the data excludes analyzing based only on a descriptor that is indicative of the parameter

or attribute of the data within the data block; identifying a first parameter or attribute of the data of the data block; compressing, if the first parameter or attribute of the data is the same as one of the plurality of parameter or attributes of the data, the data block with the at least one encoder associated with the one of the plurality of parameters or attributes of the data that is the same as the first parameter or attribute of the data to provide a compressed data block; and compressing, if the first parameter or attribute of the data is not the same as one of the plurality of parameters or attributes of the data, the data block with a default encoder to provide the compressed data block. Upon information and belief, MongoDB uses the Accused Instrumentalities, which perform the infringing method, for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to its customers.

38. MongoDB also indirectly infringes the '825 Patent by manufacturing, using, selling, offering for sale, and/or importing the accused products, with knowledge that the accused products were and are especially manufactured and/or especially adapted for use in infringing the '825 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use. On information and belief, the Accused Instrumentality is designed to function with compatible hardware to perform a method comprising: associating at least one encoder to each one of a plurality of parameters or attributes of data; analyzing data within a data block to determine whether a parameter or attribute of the data within the data block is identified for the data block; wherein the analyzing of the data within the data block to identify a parameter or attribute of the data excludes analyzing based only on a descriptor that is indicative of the parameter or attribute of the data within the data block; identifying a first parameter or attribute of the data of the

data block; compressing, if the first parameter or attribute of the data is the same as one of the plurality of parameter or attributes of the data, the data block with the at least one encoder associated with the one of the plurality of parameters or attributes of the data that is the same as the first parameter or attribute of the data to provide a compressed data block; and compressing, if the first parameter or attribute of the data is not the same as one of the plurality of parameters or attributes of the data, the data block with a default encoder to provide the compressed data block. Because the Accused Instrumentality is designed to operate as the claimed method, the Accused Instrumentality has no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. MongoDB's manufacture, use, sale, offering for sale, and/or importation of the Accused Instrumentality constitutes contributory infringement of the '825 Patent.

39. On information and belief, MongoDB has had knowledge of the '825 Patent since at least the filing of the original Complaint in this action, or shortly thereafter, and on information and belief, MongoDB knew of the '825 Patent and knew of its infringement, including by way of this lawsuit.

40. MongoDB's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe claims of the '825 Patent. Use of the Accused Instrumentalities in their ordinary and customary manner results in infringement of claims of the '825 Patent.

41. For example, MongoDB explains to customers the benefits of using the Accused Instrumentalities, such as by touting their performance advantages: "[O]nly the

differences of each successive snapshot are stored. Compression and block-level deduplication reduce the size of snapshot data.”

<https://docs.opsmanager.mongodb.com/current/application/>. For similar reasons, MongoDB also induces its customers to use the Accused Instrumentalities to infringe other claims of the '825 Patent. MongoDB specifically intended and was aware that the normal and customary use of the Accused Instrumentalities on compatible systems would infringe the '825 Patent. MongoDB performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '825 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, MongoDB engaged in such inducement to promote the sales of the Accused Instrumentalities, *e.g.*, through MongoDB's user manuals, product support, marketing materials, and training materials to actively induce the users of the accused products to infringe the '825 Patent. Accordingly, MongoDB has induced and continues to induce end users of the accused products to use the accused products in their ordinary and customary way with compatible systems to make and/or use systems infringing the '825 Patent, knowing that such use of the Accused Instrumentalities with compatible systems will result in infringement of the '825 Patent.

42. The Accused Instrumentalities associate at least one encoder to each one of a plurality of parameters or attributes of data. For example, the Accused Instrumentalities support compression and deduplication encoders (*e.g.*, “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.” <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>.) The Accused Instrumentalities analyze data blocks to detect redundant data blocks (*e.g.*, “A

process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>; “Reduces memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See* <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.) In particular, the Accused Instrumentalities associate deduplication encoder with the repeated data block and compression encoder with a unique data block.

43. The Accused Instrumentalities analyze data within a data block to determine whether a parameter or attribute of the data within the data block is identified for the data block. For example, the Accused Instrumentalities analyze data blocks to determine redundant data blocks by performing data deduplication. (e.g., “A process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>.) As another example, MongoDB storage engine analyzes index data structures to identify redundant data blocks (e.g., “Reduces memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See* <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.)

44. The Accused Instrumentalities analyze data, wherein the analyzing of the data within the data block to identify a parameter or attribute of the data excludes analyzing based only on a descriptor that is indicative of the parameter or attribute of the data within the data block. For example, the Accused Instrumentalities analyze data blocks to identify redundant data blocks by performing data deduplication (e.g., “A process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>.) As another example, MongoDB storage engine analyzes index data structures to identify redundant data blocks (e.g., “Reduces memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See* <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.) As such, the Accused Instrumentalities analyze the data blocks to detect redundant data blocks.

45. The Accused Instrumentalities identify a first parameter or attribute of the data of the data block. For example, the Accused Instrumentalities analyze the data blocks to identify redundant data blocks (e.g., “A process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>; “Reduces memory and disk consumption by storing any identical index

key prefixes only once, per page of memory.” *See*

[https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression.](https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression))

46. The Accused Instrumentalities compress, if the first parameter or attribute of the data is the same as one of the plurality of parameter or attributes of the data, the data block with the at least one encoder associated with the one of the plurality of parameters or attributes of the data that is the same as the first parameter or attribute of the data to provide a compressed data block. For example, the Accused Instrumentalities support data deduplication that compress data blocks by eliminates redundant data (e.g., “A process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See*

<https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>;

“Reduces memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See*

[https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression.](https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression))

47. The Accused Instrumentalities compress, if the first parameter or attribute of the data is not the same as one of the plurality of parameters or attributes of the data, the data block with a default encoder to provide the compressed data block. For example, the Accused Instrumentalities provide Snappy and zlib data compression techniques (e.g., “By default, collection data in WiredTiger use Snappy block compression; zlib compression is also available. *See* <https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>.) As another example, in Accused Instrumentalities, MongoDB Ops Manager

implements data compression technique during backup and archiving (e.g., “Compression and block-level deduplication reduce the size of snapshot data.” *See* <https://docs.opsmanager.mongodb.com/current/application/>.) As such, if redundant data blocks are not identified, the Accused Instrumentalities use a compression encoder to compress unique data blocks.

48. MongoDB also infringes other claims of the ’825 Patent, directly and through inducing infringement and contributory infringement.

49. On information and belief, use of the Accused Instrumentalities in their ordinary and customary fashion results in infringement of the methods claimed by the ’825 Patent.

50. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ data storage accelerating features, MongoDB has injured Realtime and is liable to Realtime for infringement of the ’825 Patent pursuant to 35 U.S.C. § 271.

51. As a result of MongoDB’s infringement of the ’825 Patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for MongoDB’s infringement, but in no event less than a reasonable royalty for the use made of the invention by MongoDB, together with interest and costs as fixed by the Court.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff Realtime respectfully requests that this Court enter:

- a. A judgment in favor of Plaintiff that MongoDB has infringed, either literally and/or under the doctrine of equivalents, the '908 Patent, the '825 Patent, and the '751 Patent;
- b. A permanent injunction prohibiting MongoDB from further acts of infringement of the '908 Patent, the '825 Patent, and the '751 Patent;
- c. A judgment and order requiring MongoDB to pay Plaintiff its damages, costs, expenses, and prejudgment and post-judgment interest for its infringement of the '908 Patent, the '825 Patent, and the '751 Patent; and
- d. A judgment and order requiring MongoDB to provide an accounting and to pay supplemental damages to Realtime, including without limitation, prejudgment and post-judgment interest;
- e. A judgment and order finding that this is an exceptional case within the meaning of 35 U.S.C. § 285 and awarding to Plaintiff its reasonable attorneys' fees against Defendants; and
- f. Any and all other relief as the Court may deem appropriate and just under the circumstances.

DEMAND FOR JURY TRIAL

Plaintiff, under Rule 38 of the Federal Rules of Civil Procedure, requests a trial by jury of any issues so triable by right.

Dated: March 12, 2019

BAYARD, P.A.

OF COUNSEL

Marc A. Fenster
Reza Mirzaie
Paul A. Kroeger
C. Jay Chung
RUSS AUGUST & KABAT
12424 Wilshire Boulevard, 12th Floor
Los Angeles, CA 90025
(310) 826-7474
mfenster@raklaw.com
rmirzaie@raklaw.com
pkroeger@raklaw.com
jchung@raklaw.com

/s/ Stephen B. Brauerman
Stephen B. Brauerman (No. 4952)
Sara E. Bussiere (No. 5725)
600 N. King Street, Suite 400
Wilmington, DE 19801
(302) 655-5000
sbrauerman@bayardlaw.com
sbussiere@bayardlaw.com

Attorneys for Plaintiff Realtime Data LLC d/b/a I XO